Welcome to ECE324

Wassily Kandinsky, *Bustling Aquarelle* (1923)

SML 310: Research Projects in Data Science, Winter 2022

Michael Guerzhoy

About me

- Michael Guerzhoy (pronounced “ger-JOY”)
- Started as Assistant Professor, Teaching Stream in EngSci last year
- Also Affiliate Scientist at St. Michael’s Hospital, working on patient chart data and data science consultant
- Previously: lecturer at Princeton, data scientist at St. Michael’s Hospital, computer vision R&D at Epson

About the class

- ECE421: basic ML + basic neural networks
- ECE368: Probabilistic models, with applications to ML
- ECE324: major project + advanced topics in ML
  - No final exam: project + presentation instead
  - No big overarching topic
    - Major topics: fairness in ML, training ML models/data science, NLP with Transformers, GANs
      - (+ whatever material we need to understand the main topics: e.g., causality, word embeddings, etc.)
      - (+whatever advanced material seems to fit in, e.g. double descent
    - Hope to run this course as more of a seminar

About the project

- Groups of 3 (one group of two allowed because of division by 3 issues; no groups of 4)
- You choose the topic
- Should be an application of ML to a domain of your choosing
- Cannot just use a standard dataset – some data collection should be done
- Topics need to be unique within the class (first come-first served)

About the class

- Initial project proposal
  - Requirements are on the website
  - Basically
    - What is the problem you are trying to address?
    - What kind of results might you expect and why? (Reference existing work)
    - Summarize the results and methods of at least two papers that addressed similar problems
  - It is understood that you may change your mind about what to work on (perhaps based on our feedback)

The first offering of ECE324 after ECE421

- Your feedback is important!